GDPS Continuous Availability

—

zExchange Technical Overview

David Aubert
GDPS tester
David.Aubert@fr.ibm.com

A few month ago, in France, a cloud data center caught fire

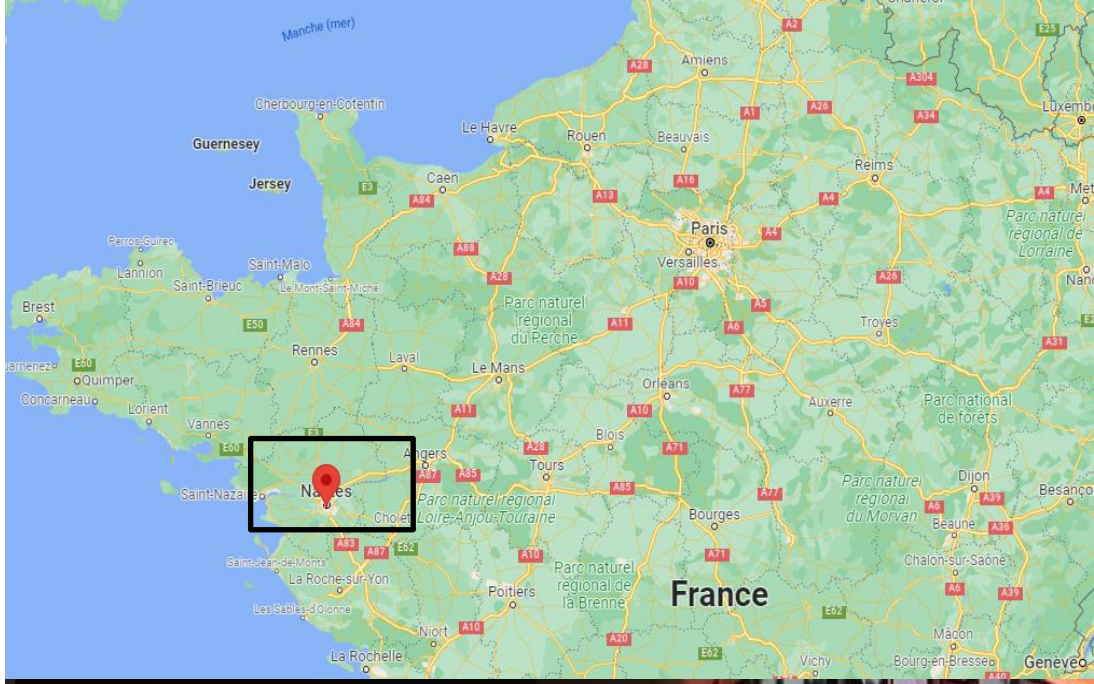

Millions of websites down
Datacenter destroyed

How much interruption can your business tolerate?

- It’s about **business continuity** and **disaster recovery**
- Saving costs and being competitive 24*7

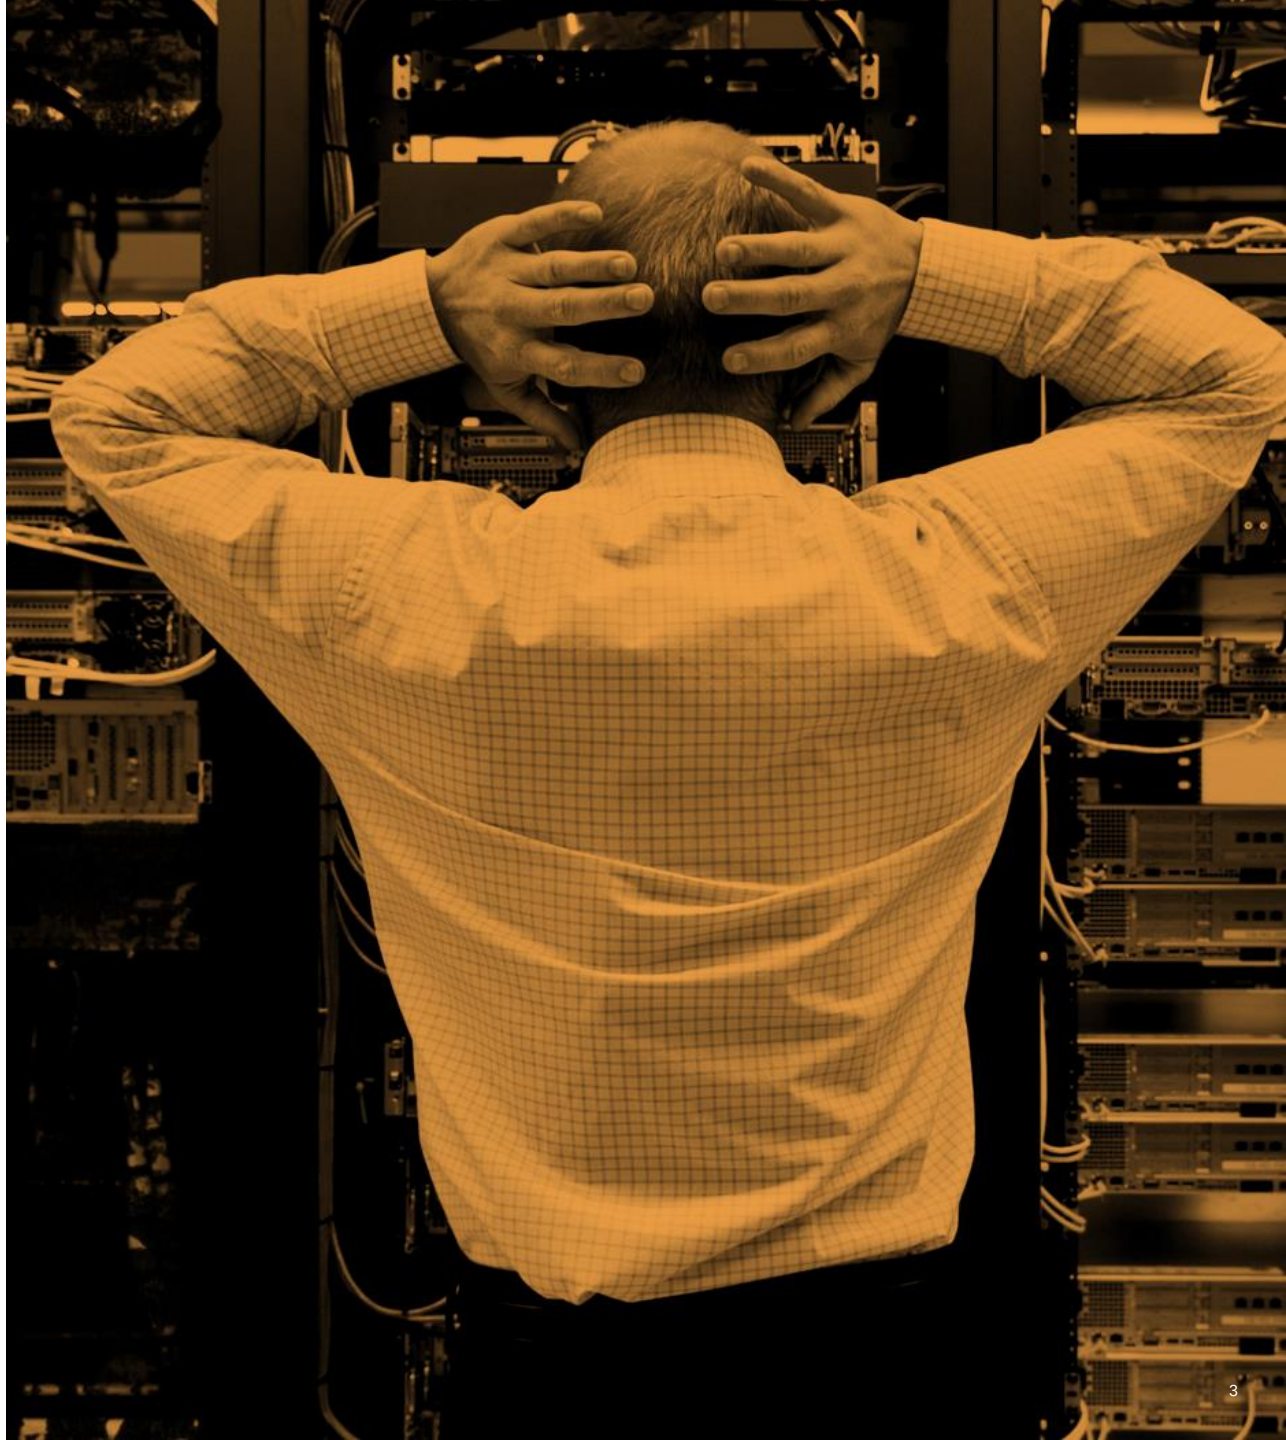

Business Continuity & Resiliency

Disruptions Remain Costly

In 2020 IBM commissioned a study of 100 IT directors in large US enterprises to understand the reality of downtime at their organization.

These IT Directors faced the following challenges:

High Cost
$5.6 M estimated cost of planned downtime in the last year.

Costs of Downtime

The average hourly costs of downtime are immense

86% of businesses lose $300,000+ per hour…

…and 34% lose $1,000,000+ per hour. *(ITIC)*

Number of Nines / Service Availability	Cost of downtime per year		
	100,000 $/**hr**	500,000 $/**hr**	1,000,000 $/**hr**
7 — 99.99999 %	$ 88	$438	$877
5 — 99.999%	$ 8,766	$43,830	$87,660
4 — 99.99%	$ 87,660	$438,300	$876,600
3 — 99.9%	$ 876,600	$4,383,000	$8,766,000
2 — 99%	$ 8,766,000	$43,830,000	$87,660,000

1. "The Real Costs Of Planned And Unplanned Downtime", Forrester Consulting, August 2019 . Forrester Opportunity Snapshot: A Custom Study Commissioned by IBM

As businesses digitally transform, IT outages are unacceptable. 40% of businesses never reopen after a disaster, of those that do, 25% fail within a year!*

Average cost per hour of downtime worldwide, by industry

- Banking / Finance avg cost per hour of down time exceeds **$9M**
- Average cost per hour of downtime across industry exceeds **$7M**

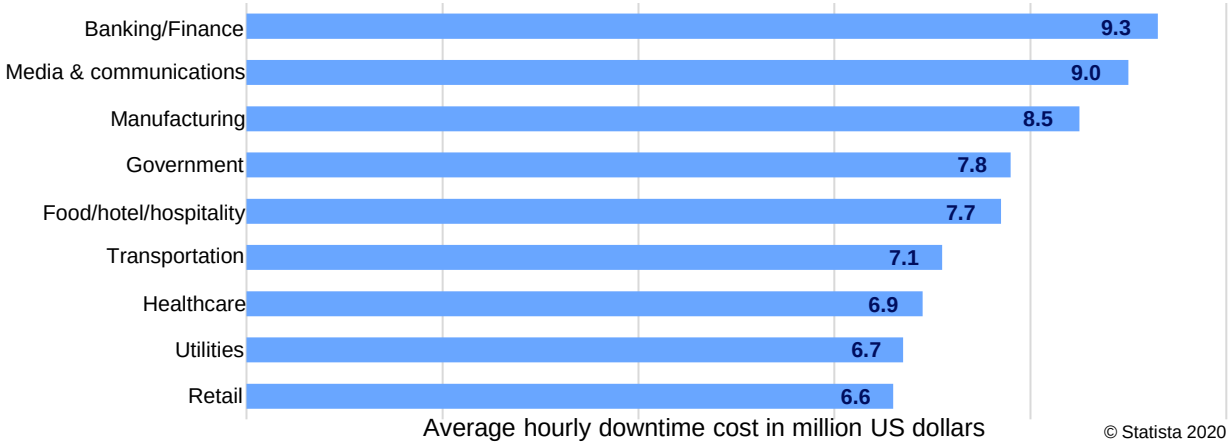

© Statista 2020

Natural Disasters

Relevant natural catastrophe loss events worldwide 2019

Tropical cyclones, extreme storms and floods caused **overall losses of US$ 150bn**

Which of the following costs does your organization face due to planned and unplanned downtime? (Forrester, 2020)

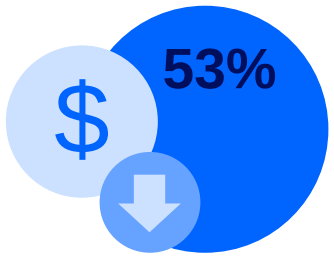

Lost revenue

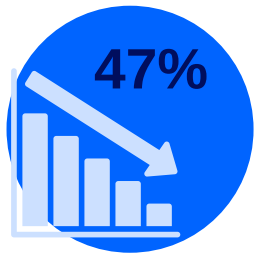

Lost productivity

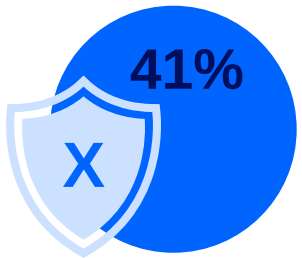

Lost brand equity or trust

Client Facing Disasters / Manmade

Problems at Facebook

Problems at Youtube

downdetector.com

Amazon explains big AWS outage, says employee error took servers offline, promises changes

BY NAT LEVY on March 2, 2017 at 9:43 am

* Access Corporation

The regulatory landscape is changing

Regulators around the globe are introducing **more stringent policies** in relation to business continuity and disaster recovery requiring more **comprehensive and extended testing** mandating clients switch over **full production loads** and operate for **30 days up to 6 months** out of their secondary data center.

FFIEC / NY DFS

Institutions should demonstrate, through testing, that their business continuity arrangements **have the ability to sustain the business until permanent operations are reestablished.**

Involve a **sufficient volume** of all types of transactions to ensure adequate capacity and functionality of the recovery facility.

Exercises generally extending over a longer period of time to **allow issues to fully evolve as they would in a crisis** and to allow realistic role-playing of all the involved groups.

EU NIS 2 Directive

EU regulators are clearly indicating the emergence of new requirements that surpass prior legislation like Operational Resiliency (ex Basel III), dealing from component failure to acknowledge risks associated to cyber attacks.

When the service is Cross-European (ex Real Time Gross Settlement,
EU Securities Settlements et cetera) EBC and EBA will supervise directly meaning **companies must adhere to a "Resiliency testing framework"**.

Regulators are asking to **prove that a secondary Site (DR) is fully functional** and can run production for a long time.

NIST Special Publication 800-53

CP-2(6) Plan for the **transfer of mission and business functions to alternate processing and/or storage sites** with minimal or no loss of operational continuity and sustain that continuity through system restoration to primary processing and/or storage sites.

CP-4(4) Include a **full recovery** and reconstitution of the system to a known state as part of contingency plan testing.

CP-7(6) Plan and prepare for circumstances that **preclude returning to the primary processing site**.

"How to protect against data loss, applications having close to 100% up time and still allow for service interrupts ?"

–

"An IT enterprise can never be down, never loose any data or allow for hostile intrusion or corruption. The market requirements becomes harder and harder each year. Yet is the enterprise obligated to keep its system on the highest possible levels of software and also to provide quick response times in a protected environment.
In short, "things has to work...every day of the year"
This session will prove that GDPS CA can become an important corner stone in the way you protect against data loss but also allow for service interruption without having your service stopped.“

–

David Aubert
GDPS tester
David.Aubert@fr.ibm.com

What is GDPS Continuous Availability (CA)?

What is the GDPS CA solution?

The GDPS CA solution provides near continuous availability for critical workloads during unplanned outages

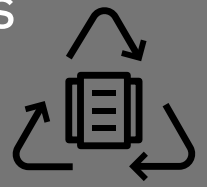

GDPS CA reduces the maintenance window for planned outages

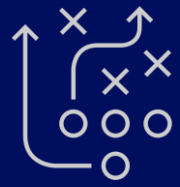

GDPS CA is a solution that enables significantly improved resiliency with data replication, workload monitoring and site switching

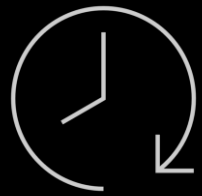

GDPS CA

Control plane to manage your business critical workloads

Automated planned/unplanned workload/site switch between sites in a matter of seconds

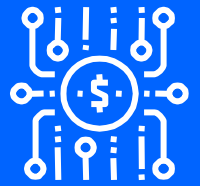

Workload monitoring

- Workload routing monitored
- Workload components monitored, STCs and JOBs
- Replication activity monitored

User Interface to manage the solution (including RESTful API)

I already have a GDPS solution! Why would I need GDPS CA?

GDPS Metro

- Disaster recovery and continuous availability
  ....BUT at limited distance

+ GDPS Continuous Availability

- Extend continuous availability
  ...to UNLIMITED distance

GDPS Global

- Disaster recovery at unlimited distance
  ...BUT no continuous availability
- Must switch whole site, impacting all business operations
- Full recovery could take over an hour

+ GDPS Continuous Availability

- Continuous availability for critical workloads
- Switch a failed workload in <30 seconds
- Keep critical workloads available during planned outages

Adding GDPS CA to an existing GDPS solution provides your operations team with a more surgical response to workload outages at unlimited distance

GDPS Continuous Availability

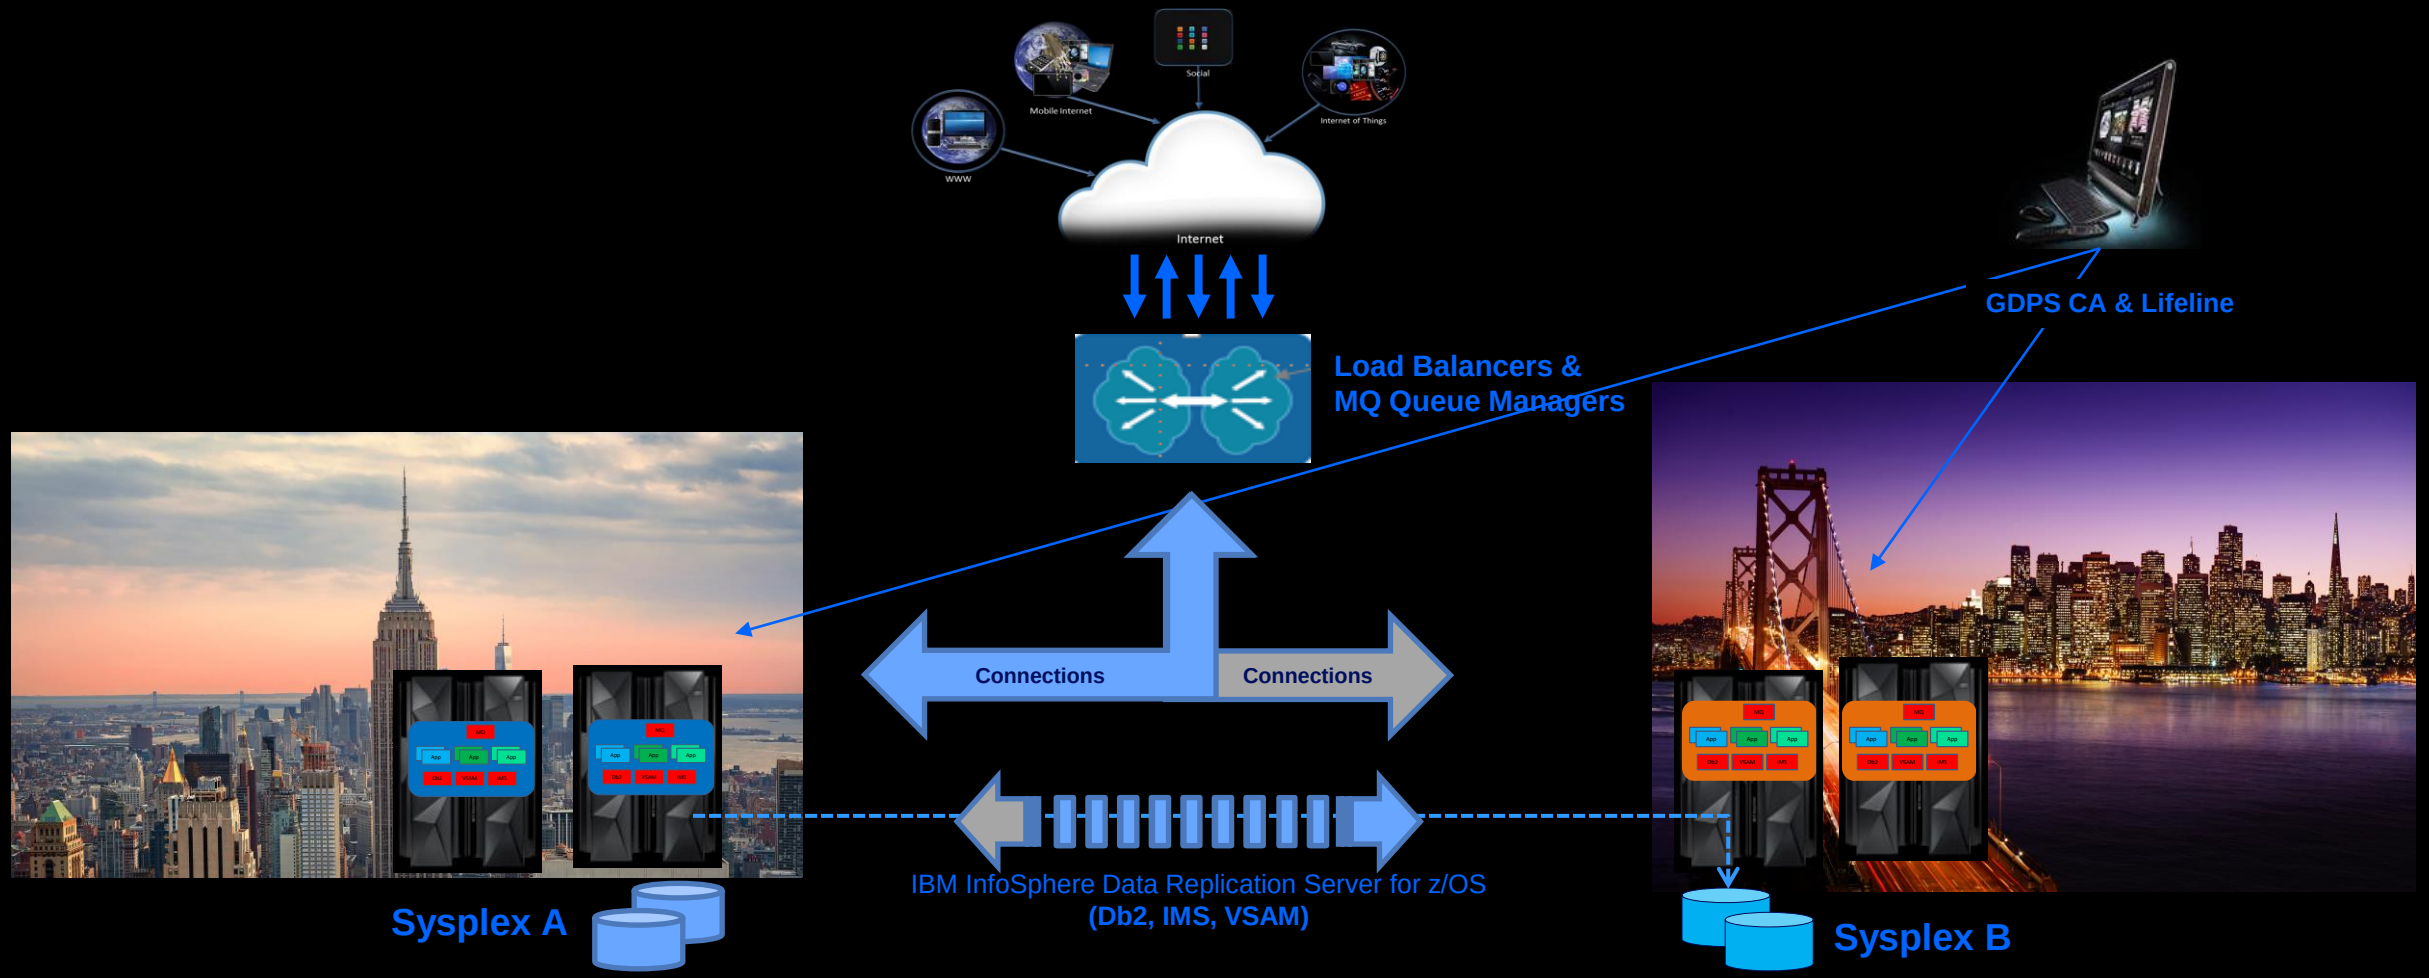

GDPS Continuous Availability Benefits

Mitigates impact
- New connections and messages routed away from failing workload applications and systems

Increases performance
- Reduces response times by routing connections and messages to workload applications and systems with capacity for additional work

Improves recovery time
- Reduces recovery time from hours (for existing DR solutions) to seconds

Unplanned workload outage

Workload elasticity
- Rerouting workloads from one site to the other with minimal disruption

Increased availability
- Outages for maintenance activities or other planned events can be minimized

Verification of disaster recovery procedures
- Non-disruptive testing of procedures by validating workloads are accessible on recovery site without requiring a production site outage

Planned workload outage

GDPS CA support for planned and unplanned outages

Distribute workload between sites

Route around failed sites

Detect workload and site failures

Perform automatic takeover or prompt for action

Switch workloads from one site to the other

Perform “graceful” takeover for planned outages

Perform “graceful” failback following a workload or site failure

Centralized monitoring of solution components

Uses software replication to maintain a transactionally consistent copy of the workload data sources on the alternate site

GDPS CA Configurations

GDPS Active/Standby Configuration

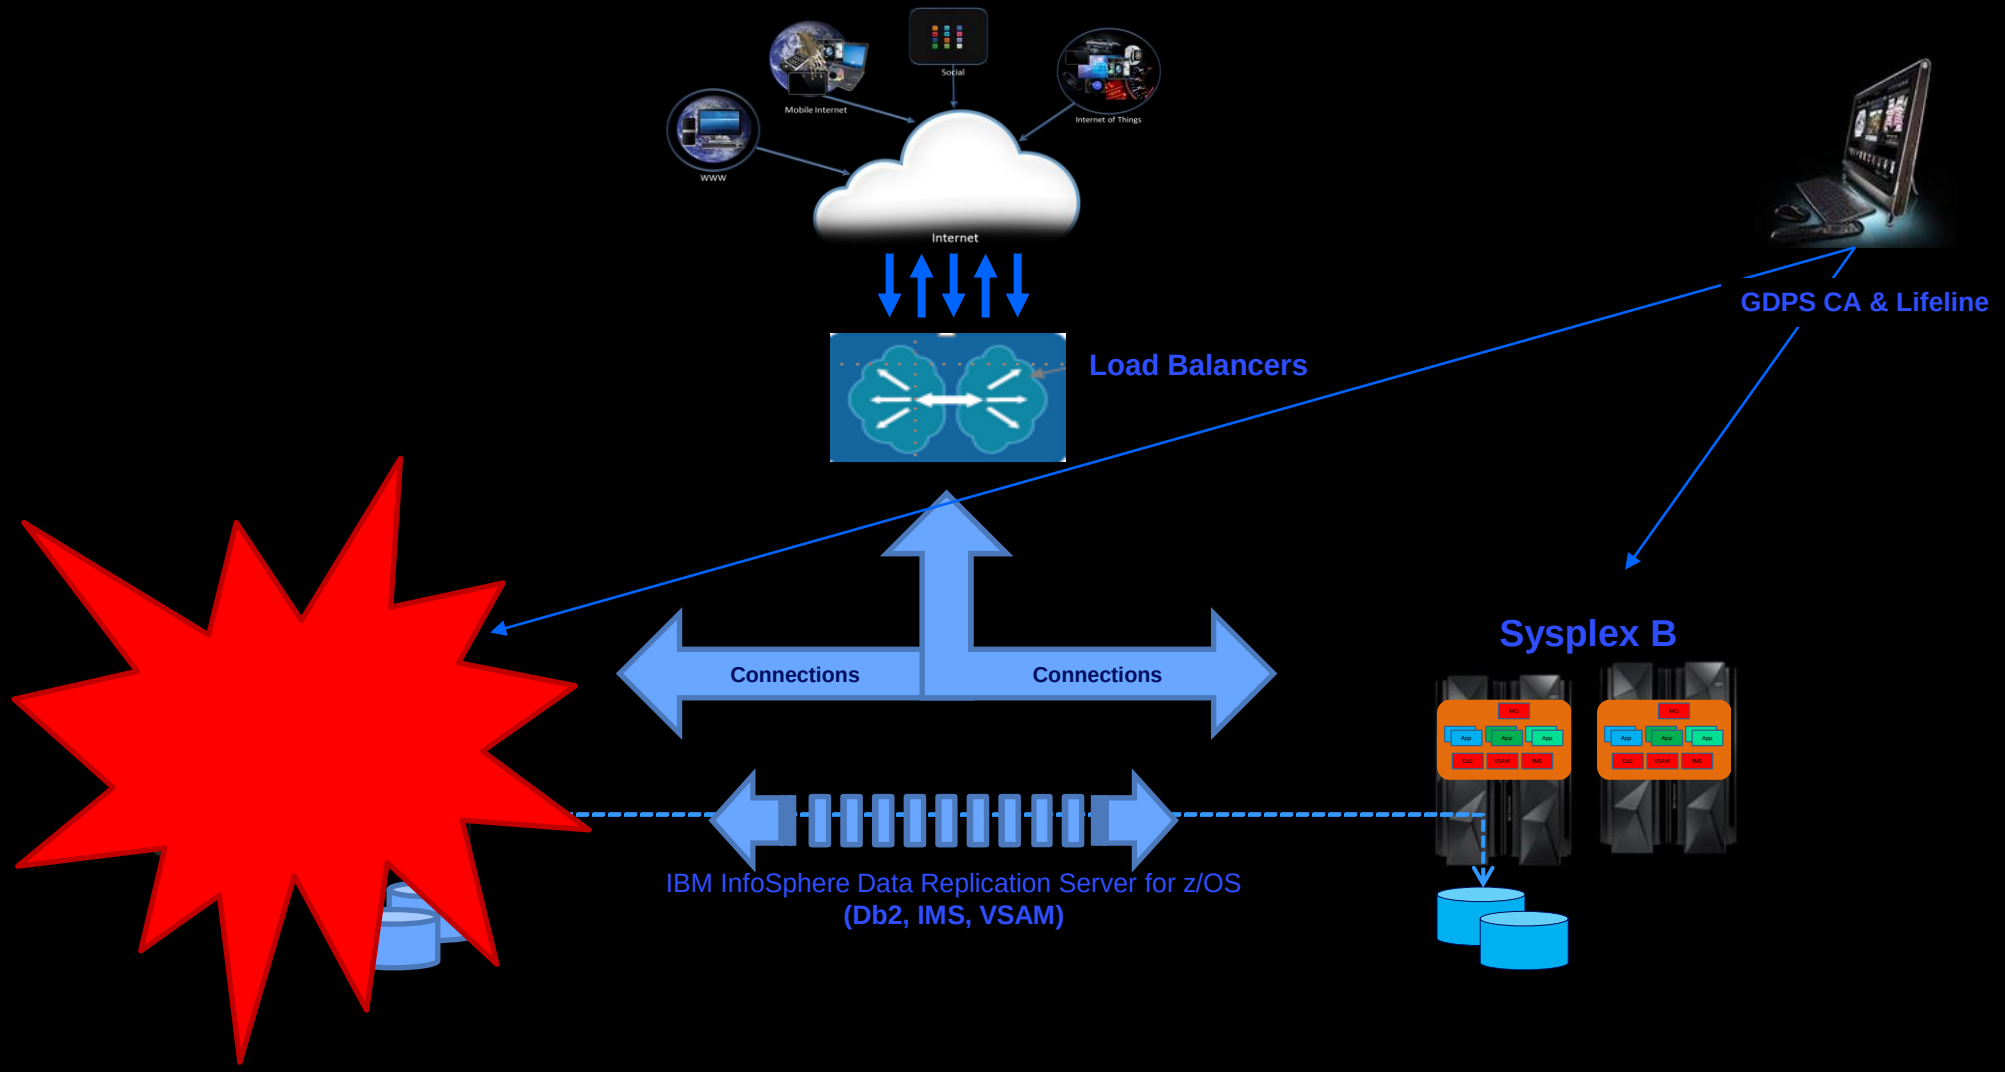

GDPS Active/Standby Configuration - MQ Cluster workloads

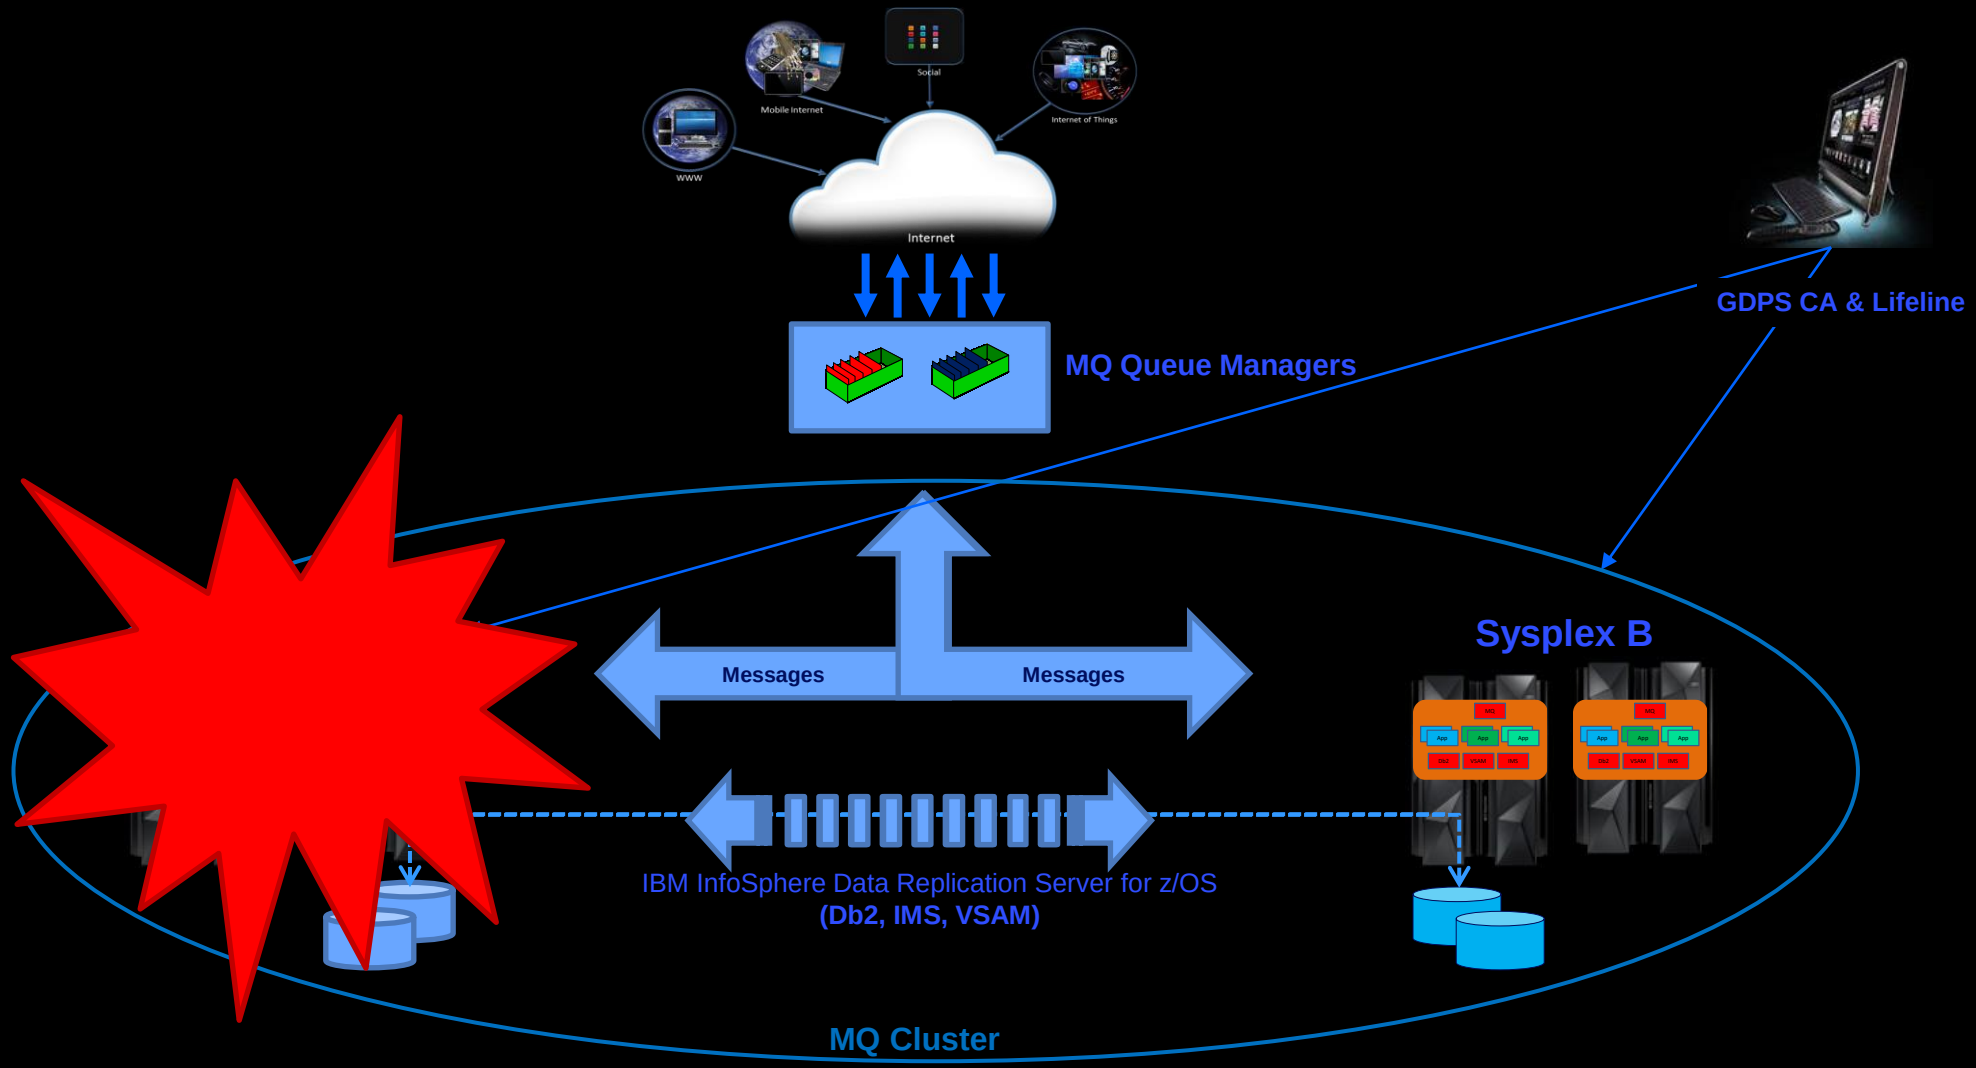

GDPS Active/Query Configuration

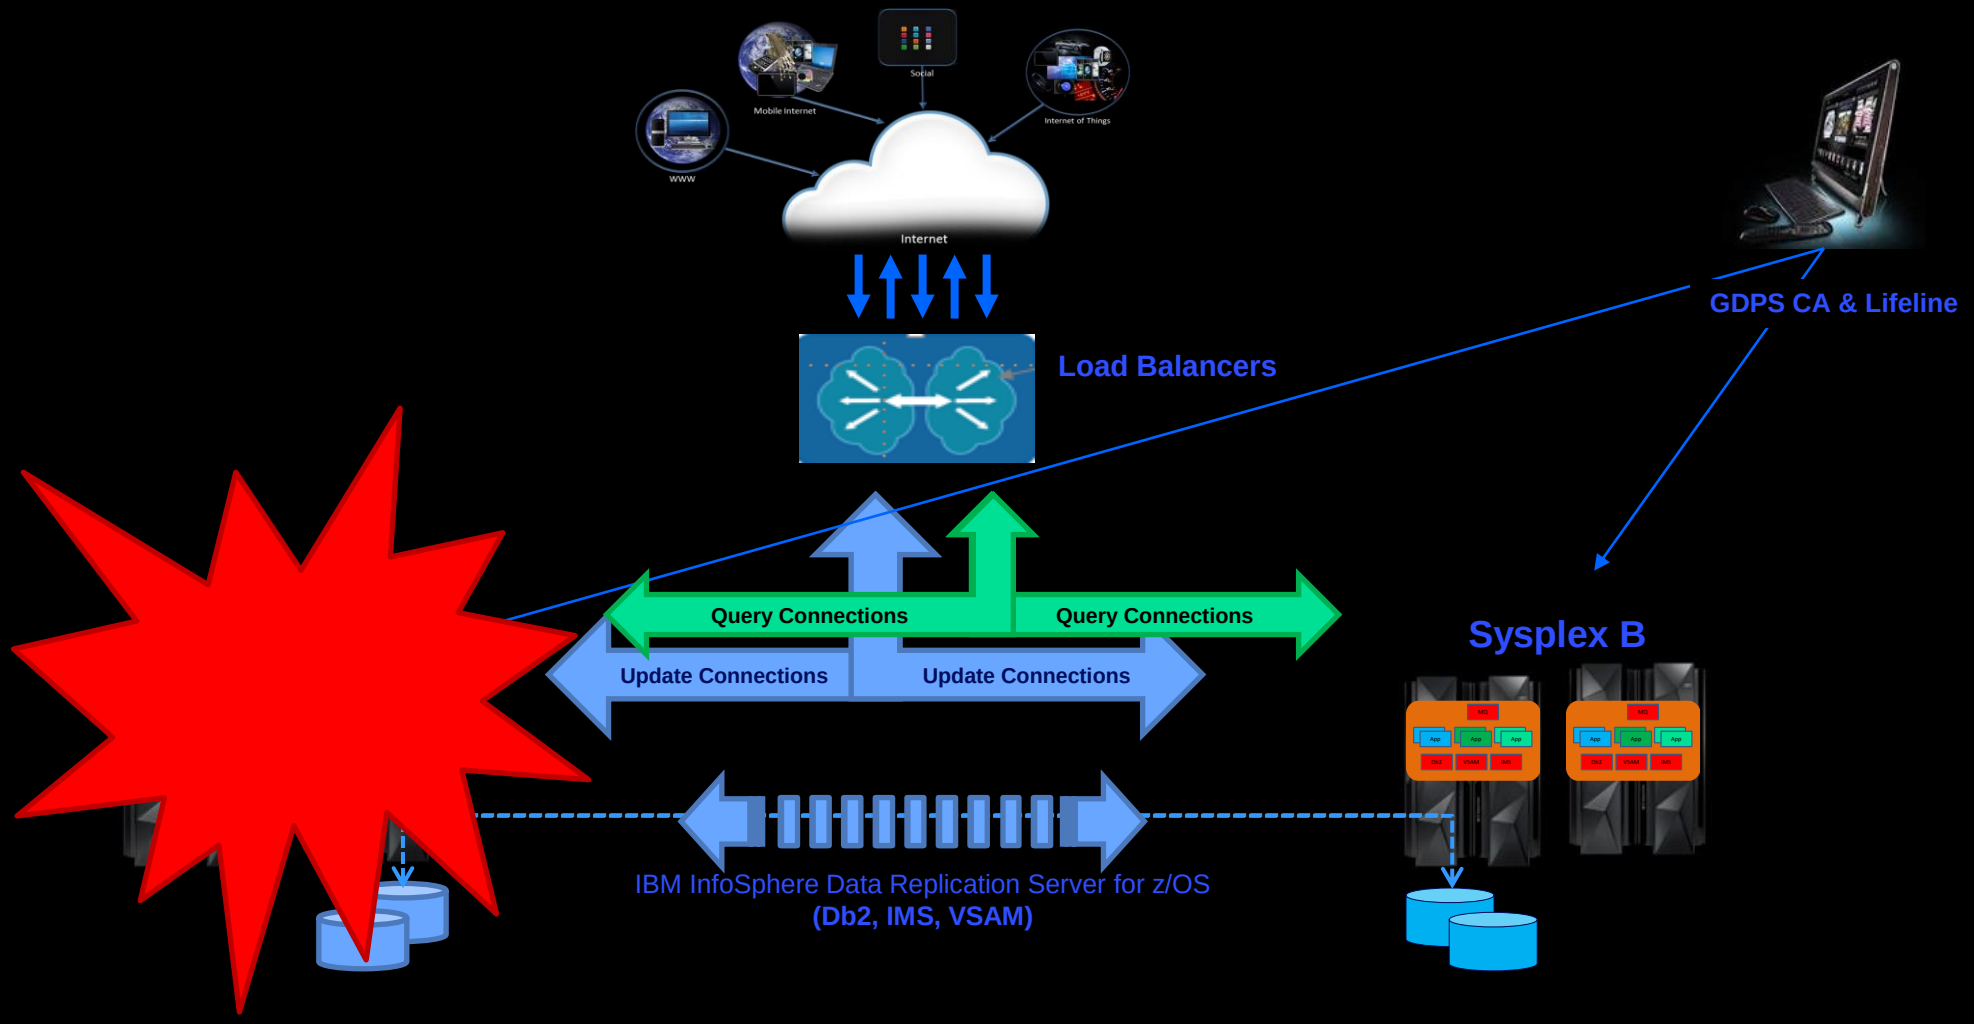

How can you use GDPS?

UI and RESTful API

Modern UI

- Using the zOS liberty server
- Web interface

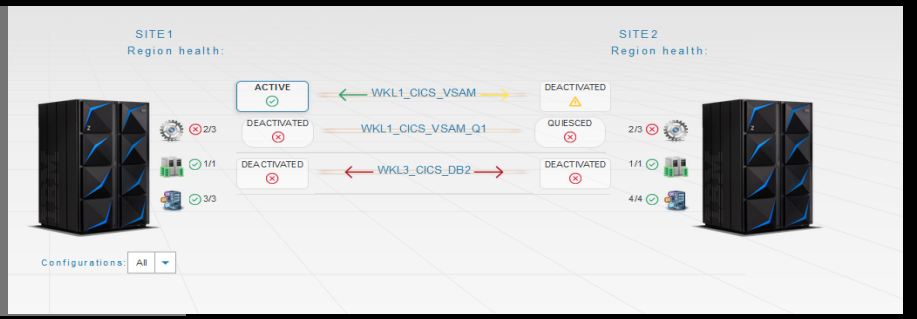

RESTful **API available**

- Web interface
- Scripts

More possibilities

- Any actions can be done using API
- You can use GDPS your own way

GDPS Continuous Availability 4.4 | Actions | Systems | TEP | Help | davida

Dashboard

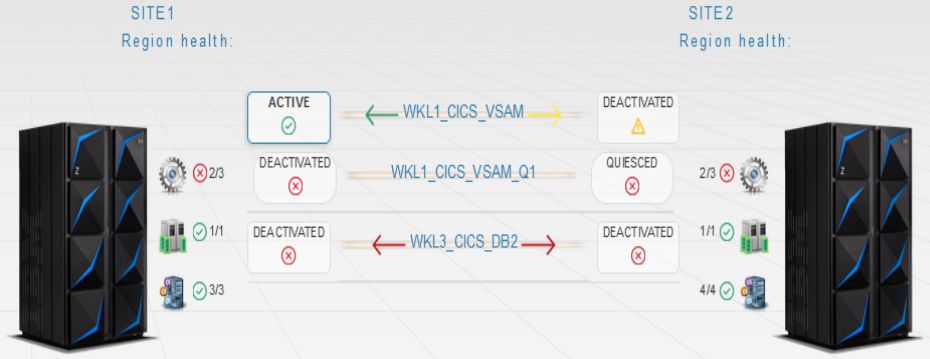

Health overview

Routing :
Replication :
Switch :

Current environment

Current System : GXC1 / NGXC1
Primary controller: GXC1 / NGXC1 (ZDL MM/GM)
Backup controller: GYC1 / (ZDL MM/GM)
GDPS version : V4.R4.M2

SDF Alerts

8 | 1
3 | 6

WTORs

1

Dashboard | WKL1_CICS_VSAM | WKL1_CICS_VSAM_Q1 | WKL3_CICS_DB2

Workload

Name: WKL3_CICS_DB2
Alias: DI
Nature / Type: DB2, IMS / Update
Max latency: 8000 ms
SPLEXXGC DB2 group: DSNDX20
SPLEXYGD DB2 group: DSNDY20
SPLEXXGC ZDL: No
SPLEXYGD ZDL: No
Site failure: PROMPT
Workload failure: PROMPT

Routing

SPLEXXGC
DEACTIVATED
UNAVAIL

SPLEXYGD
DEACTIVATED
UNAVAIL

Routing direction: None
Switch status: NO
Force switch: NO

Replication

SPLEXXGC → SPLEXYGD
DB2 consistency groups: 18 IMS subscriptions: 43

	OK	Warning	Critical
DB2 CGs	0 out of 18	0 out of 18	18 out of 18
DB2 Latency	0 out of 18	0 out of 18	18 out of 18
IMS Subscriptions	0 out of 43	0 out of 43	43 out of 43
IMS Latency	0 out of 43	0 out of 43	43 out of 43

SPLEXXGC ← SPLEXYGD
DB2 consistency groups: 18 IMS subscriptions: 43

	OK	Warning	Critical
DB2 CGs	0 out of 18	0 out of 18	18 out of 18
DB2 Latency	0 out of 18	0 out of 18	18 out of 18
IMS Subscriptions	0 out of 43	0 out of 43	43 out of 43
IMS Latency	0 out of 43	0 out of 43	43 out of 43

Workload Components (28)

SPLEXXGC: DOWN SPLEXYGD: DOWN

Sysplex	Component	System	Category	Sub category	Status
SPLEXXGC	RXC3AD11	GCP1	INFOSPHERE	DB2W311APL	INACTIVE
SPLEXXGC	RXC3AD21	GCP2	INFOSPHERE	DB2W321APL	INACTIVE
SPLEXXGC	RXC3CD11	GCP1	INFOSPHERE	DB2WK31CAP	INACTIVE
SPLEXXGC	CXC13A1	GCP1	CICS	AOR	INACTIVE
SPLEXXGC	CXC23A1	GCP2	CICS	AOR	INACTIVE
SPLEXXGC	CXC13T1	GCP1	CICS	TOR	INACTIVE
SPLEXXGC	RXCBAI01	GCP1	INFOSPHERE	IMSWB01APL	INACTIVE

Last update: 2021/11/10 08:12:22 | Refresh | Monitor last run: 2021/11/10 07:11:47 | Run monitor

Dashboard | WKL1_CICS_VSAM | WKL1_CICS_VSAM_Q1

Workload

Name: WKL1_CICS_VSAM
Alias: W1U
Nature / Type: CICSVSAM / Update
Max latency: 10000 ms
Associated Query Workloads:
WKL1_CICS_VSAM_Q1
SPLEXXGC ZDL: No
SPLEXYGD ZDL: No
Site failure: PROMPT
Workload failure: PROMPT
SPLEXXGC CICSplex: GCSPLX0
SPLEXYGD CICSplex: GDSPLX0
TSCOPE: VSAMUPD

Routing

SPLEXXGC
DEACTIVATED
AVAIL

SPLEXYGD
ACTIVE
AVAIL

Routing direction: SPLEXYGD
Switch status: YES
Force switch: YES

Replication

SPLEXXGC → SPLEXYGD

Latency threshold: 0 ms Consistency groups: 1 View all

CG group	Replication	Latency	Job status	Subscription	Job name	System
+ WKL1_CICS_VSAM	ACTIVE	0 ms	...	...	...	...

SPLEXXGC ← SPLEXYGD

Latency threshold: 573 ms Consistency groups: 1 View all

CG group	Replication	Latency	Job status	Subscription	Job name	System
+ WKL1_CICS_VSAM	ACTIVE	573 ms	...	...	...	...

Workload Components (14)

SPLEXXGC: UP SPLEXYGD: UP

Sysplex	Component	System	Category	Sub category	Status
SPLEXXGC	CXC11A1	GCP1	CICS	AOR	ACTIVE
SPLEXXGC	CXC21A1	GCP2	CICS	AOR	ACTIVE
SPLEXXGC	CXC31A1	GCP3	CICS	AOR	ACTIVE
SPLEXXGC	CXC21T1	GCP2	CICS	TOR	ACTIVE
SPLEXXGC	CXC2REP	GCP2	CICS	AOR	ACTIVE
SPLEXXGC	RXC1AV00	GCP2	INFOSPHERE	VSA0001APL	ACTIVE
SPLEXXGC	RXC1CV00	GCP2	INFOSPHERE	VSA0001CAP	ACTIVE

Last update: 2021/11/10 08:11:21 | Refresh | Monitor last run: 2021/11/10 07:12:08 | Run monitor

Health overview

Routing :
Replication :
Switch :

Current environment

Current System : GXC1 / NGXC1
Primary controller: GXC1 / NGXC1 (ZDL MM/GM)
Backup controller: GYC1 / NGYC1 (ZDL MM/GM)
GDPS version : V4.R4.M2

SDF Alerts

5 | 0
3 | 5

WTORs

1

RESTful API

Why providing RESTful API?

Large number of actions

It can be use from multiple platform

You can use multiple language

Provides access to information held in GDPS

Enables actions (Standard actions, initiate scripts, workload action etc)

RESTful API

GDPS RESTful API (Programming interface)

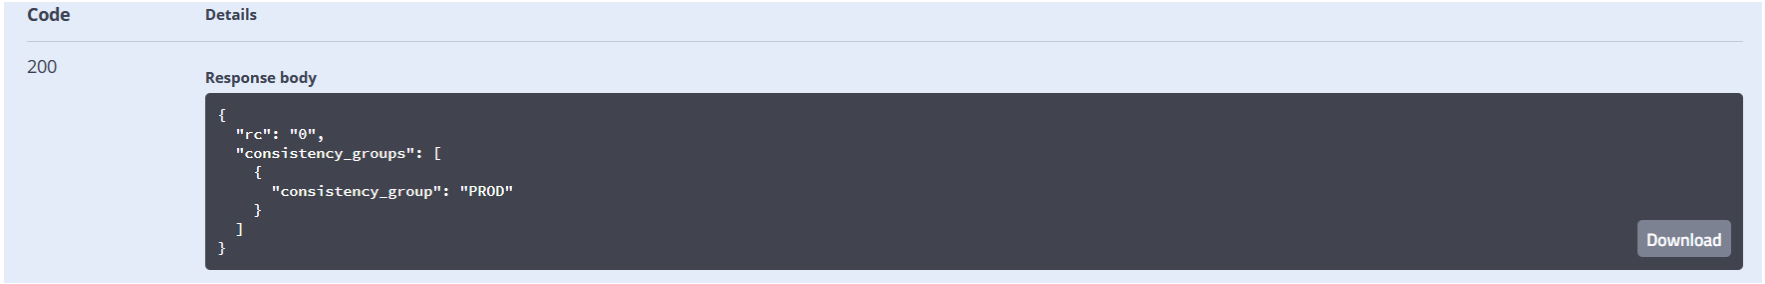

Javascript example

```
Automated_pairs_resynch {

scriptRc = run_start_secondary_script();

if(scriptRC == 0) {
    scriptEndFOund = false;
    do while(!scriptEndFound) {
        sdfTrace = get_SDF_trace_from_GDPS_API();
        display_SDF_trace_to_console();
        scriptExecStatus = get_script_execution_status(sdfTrace)
        if(scriptExecStatus.finish) {
          message = “Script STARTSEC finished with return code” + scriptExecStatus.rc +
                            “Execution trace:” + sdfTrace;
          display_to_console(message);
          send_email_to_operation(message);
        }
  }
}
```

Planned switch versus Unplanned

Why is there a difference?

Remember:

"Detect workload and site failures

Perform automatic takeover or prompt for action"

For more flexibility you can create scripts that will be triggered in case of unplanned switch to fit the requisite of *YOUR* environment

The unplanned switch is automatically detected but *YOU* choose how to answer to the failure

Multi-site Workload Lifeline

Multi-site Workload Lifeline workload monitoring

Determine whether workload applications are active, healthy, and meeting their SLA objectives

Verify resources and status of systems and sites where workload applications reside

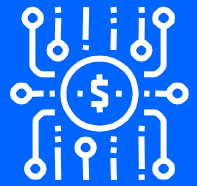

Monitor different types of workloads

- TCP applications
- SNA appls
- MQ cluster applications
- Db2 Sysplex Routed applications
- Applications on Linux on Z

React quickly (within seconds) to changes in the workload to influence how workload connections and MQ messages are routed

Multi-site Workload Lifeline workload routing

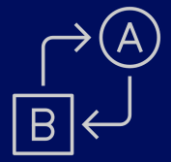

Adjust workload routing recommendations quickly in response to changes in workload availability

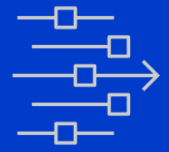

Provide controlled rerouting of workloads for planned outages

Support different routing configurations based on workload requirements

- Active/Standby
- Active/Query

Single point of control to influence workload connection routing and workload MQ message routing

Application Servers and Databases

Application Servers and Databases

- Workloads managed by GDPS CA include those run in CICS TS, IMS TM and batch
- Eg CICS-Db2 , IMS-Db2, IMS-DB, CICS-VSAM apps

Db2 and IMS provide replication logging to record database updates

CICS TS provides replication logging for online VSAM updates

CICS VR provides replication logging for batch VSAM updates

Customer case study

Leading Bank Realizes over 99% Improvement with GDPS Continuous Availability

One of the world's largest banks wanted to

- Reduce their planned outage downtime
- Avoid unplanned outages for critical workloads
- Better distribute production load across data centers

With GDPS CA solution

Reduced planned downtime from 4 hours to under 30 seconds

What is the GDPS CA solution?

The 3 items to remember:

1. It improves your resiliency
2. Disaster recovery AND high availability
3. Protect you from data loss

The GDPS CA solution provides near continuous availability for critical workloads during unplanned outages

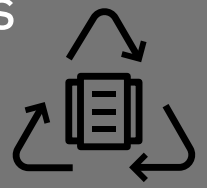

GDPS CA reduces the maintenance window for planned outages

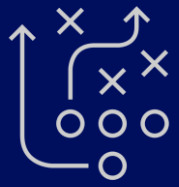

GDPS CA is a solution that enables significantly improved resiliency with data replication, workload monitoring and site switching

Questions?

david.aubert@fr.ibm.com

Resources

- IBM Product page for GDPS:
  https://www.ibm.com/it-infrastructure/z/technologies/gdps
- IBM Product page for Lifeline:
  https://www.ibm.com/products/multisite-workload-lifeline
- IBM Documentation for Lifeline:
  https://www.ibm.com/docs/en/mwl
- IBM Documentation for IBM Z NetView for Continuous Availability
  https://www.ibm.com/docs/en/znfca/6.3.0
- IBM Documentation for IIDR for Db2
  https://www.ibm.com/docs/en/idr/11.4.0?topic=overviews-q-replication-information-roadmap#iiyrqinfroadmap
- IBM Documentation for IIDR for IMS
  https://www.ibm.com/docs/en/iirfz/11.3.0
- IBM Documentation for IIDR for VSAM
  https://www.ibm.com/docs/en/idrfvfz/11.3.0

Backup - Components

IBM InfoSphere Data Replication

IBM InfoSphere Data Replication

Captures change records from source DBMS logs

Replicates transactions by workload, forwarding only committed changes.

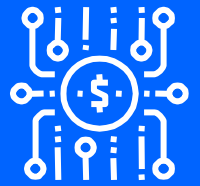

Achieves near real-time synchronization, by transmitting changed data across multiple channels and using multiple target DBMS connections for parallel transaction replay.

Fully controlled from the GDPS CA Control Plane, with centralized monitoring and operations.

Z NetView for Continuous Availability

Z NetView for Continuous Availability

Provide functional enablement for GDPS Continuous Availability components in NetView address space

Process replication events containing latency information that influence workload and site switches

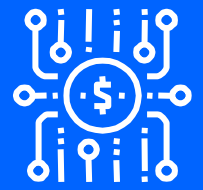

Monitor key performance metrics

- Multi-site Workload Lifeline
- IIDR for Db2
- IIDR for IMS
- IIDR for VSAM

Centralized location for event processing, monitoring information, and communication with NetView enterprise master technology

GDPS Continuous Availability

"How to protect against data loss, applications having close to 100% up time and still allow for service interrupts ?"

Abstract:
"An IT enterprise can never be down, never loose any data or allow for hostile intrusion or corruption. The market requirements becomes harder and harder each year. Yet is the enterprise obligated
to keep it's system on the highest possible levels of software and also to provide quick response times in a protected environment.
In short, "things has to work...every day of the year"
This session will prove that GDPS CA can become an important corner stone in the way you protect against data loss but also allow for service interruption without having your service stopped."